

WE CLAIM:

1. An interconnection controller, comprising:

an intra-cluster interface configured for coupling with intra-cluster links to a plurality of local processors arranged in a point-to-point architecture in a local cluster;

an inter-cluster interface configured for coupling with an inter-cluster link to a non-local interconnection controller in a non-local cluster;

a transceiver configured to receive an intra-cluster packet from the local processors via the intra-cluster links of the coherent interface and encapsulate the intra-cluster packet as a high-speed link packet by adding a sequence identifier and a packet type identifier to the intra-cluster packet; and

a serializer/deserializer configured to serialize the high-speed link packet and forward the serialized high-speed link packet to the inter-cluster interface for transmission to the non-local interconnection controller via an inter-cluster link.

2. The interconnection controller of claim 1, further configured to compute a cyclic redundancy code value based only on the high-speed link packet, and wherein the serializer/deserializer is further configured to encode the cyclic redundancy code value in the high-speed link packet.

3. The interconnection controller of claim 1, wherein the inter-cluster interface is further configured to receive high-speed link packets from the non-local interconnection controller, wherein the serializer/deserializer is further configured to deserialize the encoded, serialized packets and wherein the transceiver is further configured to perform a cyclic redundancy code check on the deserialized packets.

4. The interconnection controller of claim 1, wherein the sequence identifier and the packet type identifier are encoded in a header of the high-speed link packet.

5. The interconnection controller of claim 1, wherein the sequence identifier and the packet type identifier are encoded in a portion of the high-speed link packet reserved for link-layer encoding.

6. An integrated circuit comprising the interconnection controller of claim 1.

7. A set of semiconductor processing masks representative of at least a portion of the interconnection controller of claim 1.

8. At least one computer-readable medium having data structures stored therein representative of the interconnection controller of claim 1.

9. The interconnection controller of claim 1, wherein the transceiver is further configured to encode a packet length field in a header of the high-speed link packet.

10. The interconnection controller of claim 1, wherein the transceiver is further configured to encapsulate both an intra-cluster packet and an ACK packet in a single high-speed link packet.

11. The interconnection controller of claim 1, wherein the transceiver is further configured to form high-speed link packets having lengths that are integral multiples of 96 bits.

12. The interconnection controller of claim 1, wherein the transceiver is further configured to form high-speed link packets that are transparent to a protocol layer.

5 13. The integrated circuit of claim 6, wherein the integrated circuit comprises an application-specific integrated circuit.

14. The at least one computer-readable medium of claim 8, wherein the data structures comprise a simulatable representation of the interconnection controller.

10 15. The at least one computer-readable medium of claim 8, wherein the data structures comprise a code description of the interconnection controller.

16. The at least one computer-readable medium of claim 14, wherein the
15 simulatable representation comprises a netlist.

17. The at least one computer-readable medium of claim 15, wherein the code description corresponds to a hardware description language.

20 18. A computer system comprising a plurality of processor clusters interconnected by a plurality of point-to-point inter-cluster links, each processor cluster comprising nodes including a plurality of local processors and an interconnection controller interconnected by a plurality of point-to-point intra-cluster links, communications within a cluster being made via an intra-cluster protocol that uses intra-cluster packets, wherein the
25 interconnection controller in each cluster is operable to map locally-generated communications directed to others of the clusters to the point-to-point inter-cluster links and

to map remotely-generated communications directed to the local nodes to the point-to-point intra-cluster links, communications between clusters being made via an inter-cluster protocol that uses inter-cluster packets, an inter-cluster packet encapsulating at least one intra-cluster packet, each interconnection controller configured to:

- 5 receive an intra-cluster packet from the local processors via the intra-cluster links of the coherent interface;
- encapsulate the intra-cluster packet as a high-speed link packet by adding a sequence identifier and a packet type identifier to the intra-cluster packet;
- serialize the high-speed link packet; and
- 10 forward the serialized high-speed link packet to the inter-cluster interface for transmission to another processor cluster via an inter-cluster link.